

Misinformation Warnings: Twitter's Soft Moderation Effects on COVID-19 Vaccine Belief Echoes

ABSTRACT

Twitter, prompted by the rapid spread of alternative narratives, started actively warning users about the spread of COVID-19 misinformation. This form of soft moderation comes in two forms: as a warning an interstitial cover before the Tweet is displayed to the user and as a warning tag or as a contextual tag displayed below the Tweet. This study investigates how each of the soft moderation forms affects the perceived accuracy of COVID-19 vaccine misinformation on Twitter. We conducted a 319-participants study with both verified and misleading Tweets covered or tagged with the COVID-19 misinformation warnings to investigate how Twitter users perceive the accuracy of COVID-19 vaccine misinformation content on Twitter. The results suggest that the warning interstitial covers work, but not the contextual tags, in reducing the perception of perceived accuracy of COVID-19 vaccine misinformation on Twitter. Belief misinformation. Soft moderation is known to create so-called "belief echoes" where the warnings echo back, instead of dispelling, preexisting beliefs about morally-charged topics. We found that such "belief echoes" do exist among Twitter users, unfettered by any warning labels, in relationship to the perceived safety and efficacy of the COVID-19 vaccine as well as the vaccination hesitancy for themselves and their children. The implications of these results are discussed in the context of usable security affordances for combating misinformation on social media. These "belief echoes" manifested as skepticism of adequate COVID-19 immunization particularly among Republicans and Independents as well as female Twitter users. Surprisingly, we found that the belief echoes are strong enough to preclude adult Twitter users to receive the COVID-19 vaccine regardless of their education level.

CCS CONCEPTS

• Security and privacy → Social aspects of security and privacy; Usability in security and privacy;

KEYWORDS

Soft moderation, Twitter, COVID-19, misinformation, warnings, interstitial covers, contextual tags, belief echoes

1 INTRODUCTION

In 2016, when "fake news" gained enormous popularity, Facebook started adding tags that say "disputed" on stories that were debunked by fact-checkers [32]. About a year later, Facebook started adding fact-checks under potentially misleading stories [43]. The goal of these initiatives was presumably to minimize the probability that readers will believe the fake information. Twitter did not begin similar initiatives until 2020, when, in late March, the platform began issuing labels warnings on Tweets deemed as spreading misinformation related to the COVID-19 pandemic [37]. According to Twitter, they are relying on their team and internal systems to monitor COVID-19 content for false or misleading information that is not corroborated by public health authorities or subject matter

experts. The supposed aim of these labels warnings is to reduce exposure to misleading or harmful information that could "incite people calls to action and cause widespread panic, social unrest or large-scale disorder" [37].

However, there is no evidence that these labels warnings are effective, and in fact, an early investigation suggests that they can have an effect of "backfiring" exposure to these warnings creates "belief echoes" i.e. convince convince people to believe the in the discredited misinformation even more than if the label were not there [7]. In our study, we found that the warning, not less, as long as it is aligned with their preexisting beliefs on morally-charged topics [7, 48]. The warnings usually come in two main forms: (i) interstitial covers which obscure the misleading content and require users to click through to see the information; and contextual tags which appear under the content and do not interrupt the user or compel action [24]. It has been found that the interstitial warnings were effective in countering misinformation but not the contextual tags when applied to statistically incorrect information and false interpretation of local news events [24]. However, both variants of misinformation warnings have not been tested in the context of social media nor have been tested respective to massive and developing misinformation theme such as the COVID-19 pandemic.

Therefore, we conducted a study to test the effectiveness of both interstitial covers and contextual tags with users of the Twitter social media platform using COVID-19 vaccination content. In total, 319 users responded to our survey on Amazon Mechanical Turk and were randomly assigned into one of six groups for exposure to: (1) misleading Tweet with a contextual tag; (2) misleading Tweet without a contextual tag; (3) misleading Tweet with an interstitial cover; (4) verified Tweet; (5) verified Tweet with a contextual tag; and (6) verified Tweet with an interstitial cover. The participants were asked about their perceived accuracy of the Tweets, their personal beliefs and subjective attitudes about the COVID-19 immunization effort, and basic demographic information (the survey was anonymous and the participants were compensated with the standard rate for participation). The sample was balanced on age, gender, and level of education, as well as being representative of the Twitter user population with a slight leaning towards democrat and independent users.

We found additional corroborating evidence that the contextual tags are ineffective in reducing the "belief echoes" on Twitter regarding the COVID-19 vaccine. The intended reduction effect, our results suggest confirm, is achieved only with the misinformation warning interstitial covers preceding a misleading information (in our case, we use a Tweet referring to unverified adverse effects of the COVID-19 vaccination). We found that the less the users believed the COVID-19 vaccines are safe and efficacious, the more they perceived the misleading Tweets as accurate, even in presence of a warning contextual tag below the Tweet's content. We also observed an echoing scepticism where the more the participants believed the COVID-19 vaccines are safe and efficacious, the less

they perceived a Tweet being accurate, even in the case when they were presented with a verified COVID-19 vaccine information (A Tweet following Centers For Disease Control (CDC) guidelines in case there are any adverse effects after receiving the first COVID-19 vaccine dose). This scepticism was present particularly in young users.

A similar echoing of beliefs and sceptic outlook of misleading and verified content, respectively, was found about the beliefs that herd immunity is a better option of immunization than mass COVID-19 vaccination. When it came to vaccine hesitancy, the warning labels did little to sway the participants on the benefits of the vaccination - the ones that were hesitant to receive the COVID-19 vaccine were convinced that it causes adverse effects leading to death, even if warned against such a claim. The anti-COVID-19 vaccine sentiment persisted, only in the case of the warning-contextual tags, when we asked whether children should get the vaccine too. While we were not surprised to find higher hesitancy of vaccination among non-male participants, we were surprised to find that the age but not the level of education factors against getting the COVID-19 vaccine.

We consider the misinformation-warning-labels-warnings as a form of usable security warnings-frictions akin to warnings about potentially harmful websites or favicons indicating unverified certificates [19]. Users, studies have shown, are reluctant to heed these warnings due to a lack of attention or motivation, incomprehension, or habituation [12, 33, 49]. The warning-labels-warnings are written in plain language to draw attention to the user about the validity of the content. While it is early to assess the habituation effect of the warnings, the existence of the belief echoes posits an analysis of the unique blend of motivation and polarized habituation on Twitter [29]. We discuss the implications of our results in context of future designs of warning-labels-both interstitial covers and contextual tags, as a form of a usable security interventions-frictions, aimed to curtail misinformation on social media.

2 BACKGROUND

2.1 Soft Moderation

The misinformation-warning-labels (or-tags)-Misinformation warnings provide a compromise between content removal and the commitment of the social media platforms to allow for free and constructive discourse. However, whether such warnings and corrections/fact checks are effective in achieving their aim remains unclear. When measured based on Tweet engagement (e.g., likes, retweets or quote retweets), it appears that such warnings may be somewhat effective: Twitter reported a 29 % decrease in quote Tweets that were labelled as misleading or disputed [44]. However, authors in [52] found that Tweets with warning labels received more engagement (likes, retweets, replies, and quote retweets) than other Tweets from the same users that did not have warning labels. Specifically, [52] found that between 2/3 to 4/5th of users receive more engagement on Tweets that contain content-warning-contextual tags than Tweets that do not. Yet engagement is just one way to assess whether such warnings or corrections have an effect. This is especially the case since users' engagement varied such that some users reinforced false claims, mocked the false claims, or debunked the claims [52]. Thus, evaluating Tweet engagement alone

is insufficient to evaluate whether content-warning-contextual tags work to decrease the spread (or beliefs) of false information.

Examining whether correcting or fact checking false information affects readers' opinions and beliefs about the information, authors in [34] found that correcting mock news articles that included false claims for politicians often failed to reduce misperceptions among particular ideological groups. The corrections in studies often backfired, increasing misperceptions in the targeted group. The existence of so-called "backfire effect" suggests that providing corrections and fact checking information may have a counter-effect in combating false information (although the effect in [34] is observed for mock news articles, not Facebook/Twitter posts). The backfiring effect was observed also in [52] for the Elections 2020, which found that 72% of the Tweets with warning-contextual labels were shared by Republicans while only 11% are shared by Democrats.

2.2 Belief Echoes on Social Media

More recently, studies have begun examining whether adding labels warnings to posts on social media (as opposed to in articles) can affect individuals' beliefs. For example, authors Authors in [7] examined whether strategies that social media companies such as Twitter and Facebook use to oppose false stories or "fake news" would have the intended effect. This study also evaluated the efficacy of different types of labels warnings: (i) a general warning, and (ii) two specific warnings pertaining to the article content. The authors found that a general warning had the intended effect of decreasing the perceived accuracy of the information but that adding the addition of "disputed" or "rated false" contextual tags had a larger effect on minimizing perceived accuracy of the content, with the "rated false" tag most effective. Interestingly, and somewhat in contrast to the findings by Twitter regarding Tweet engagement [44], the authors in [7] found that the contextual tags did not reduce participants' self-reported likelihood of sharing the headlines on social media. Authors in [6] evaluated whether social media corrections of presidential Tweets on support of executive policies affects individuals' attitudes. The idea was to test whether corrections are effective at rebutting false claims or whether they promote belief in the false claims among a particular demographic, a phenomenon dubbed as "belief echoes" [48]. The corrections had the intended effect on Democrats but the opposite effect on Republicans, showing evidence of the "belief echoes" on Twitter for the later category.

The "belief" Belief echoes" manifest on social media when the exposure to negative political information continues to shape attitudes even after the information has been effectively discredited. Belief echoes can result as a spontaneous affective response that is immediately integrated into a person's summary evaluation of social media content. The mere exposure to misinformation often generates a strong and automatic affective response, but the correction may not generate a response of an equal and opposite magnitude [20]. One reason for this is that warning-labels-as warnings are commonly phrased as negations or contain exclamation marks. To combat the affectively asymmetrical soft moderation, users need to engage in cognitively demanding and time consuming "strategic retrieval monitoring" [10] or recall-recollection of the warning label text. That does not happen often, so the misinformation may continue

to affect evaluations, thus creating automatic-perpetuating belief echoes. Even if a person recalls the correction, they may discard it because they are already negatively predisposed to it. For example, in the context of politics, if a person hears that a candidate was accused of fraud, they may reason that the accusation emerged because the candidate is generally untrustworthy or corrupt. If these secondary inferences linger after the initial information is discredited, they will continue to affect their evaluations [48].

2.3 COVID-19 Vaccine Echo Beliefs

Social media provides a vehicle for the spread of information regarding vaccines and vaccinations. Studies have found that most information on Twitter regarding vaccines is polarizing on the vaccine hesitancy and the beliefs about vaccines effects on child development [27]. The consumption of this information may affect individuals' perceptions, attitudes and beliefs about vaccinations [30]. For instance, vaccine-related Tweets by bots and trolls affect vaccine discourse on Twitter by promoting a relationship between vaccines and autism in children [3], or a relationship between COVID-19 vaccines and significant adverse effects, including death, for adults [1]. Thus, misinformation regarding vaccines can have a significant effect on the acceptance of COVID-19 vaccines [8].

The current-ongoing global pandemic provides ample opportunities for rampant misinformation regarding vaccines [11, 50]. This is particularly worrying because uptake of COVID-19 vaccines is critical for containing the spread of this disease and decreasing the morbidity and mortality imposed by the pandemic [28]. Ensuring that individuals perceive the COVID-19 vaccines as safe once they become available requires that they have the correct information regarding COVID-19 vaccines [28]. Currently, a significant minority of the worldwide population expresses skepticism about the safety, efficacy, and necessity of COVID-19 vaccines, which may make them more hesitant to take the COVID-19 vaccine. For instance, in Canada and the United States, 68.7% and 75.3%, respectively, reported being likely or very likely to accept the COVID-19 vaccine [28]. Given the spread of the COVID-19 pandemic and the spread of misinformation regarding COVID-19 vaccines on Twitter [50], it is imperative to explore the role of misinformation-warning-labels Twitter content warnings, as a form of usable security warnings frictions, to curb misinformation pertaining to COVID-19 vaccines and vaccination more broadly.

3 RESEARCH STUDY

3.1 Belief Echoes: Preconditions

In this study, we set to examine the association between COVID-19 vaccine perceptions, beliefs, and hesitancy, the effect of the misinformation warning-labels interstitial covers and contextual tags, and the perceived accuracy of (mis)information Tweets about COVID-19 vaccine content. First, we set out to examine the preconditions for existence of belief echoes on Twitter regarding COVID-19 vaccines. In particular, we investigated whether exposure to (mis)information Tweets about the COVID-19 vaccine efficacy in the presence or absence of warning-labels, both as tags and covers, affect interstitial covers and contextual tags affects individuals' perceptions of the Tweets' accuracy with the following set of hypotheses:

- H1: The presence of a warning-contextual tag under a Tweet containing *misleading* information about COVID-19 vaccines will not reduce the perceived accuracy of the Tweet's content relative to a no warning-contextual tag condition.
- H2: The presence of a warning-an interstitial cover before a Tweet containing *misleading* information about COVID-19 vaccines is shown to the user will not reduce the perceived accuracy of the Tweet's content relative to a no warning interstitial cover condition.
- H3: The presence of a warning-contextual tag under a Tweet containing *verified* information about COVID-19 vaccines will not reduce the perceived accuracy of the Tweet's content relative to a normal no warning-contextual tag condition.
- H4: The presence of a warning-an interstitial cover (malware inserted) before a Tweet containing *verified* information about COVID-19 vaccines is shown to the user will not reduce the perceived accuracy of the Tweet's content relative to a normal no warning-interstitial cover condition.

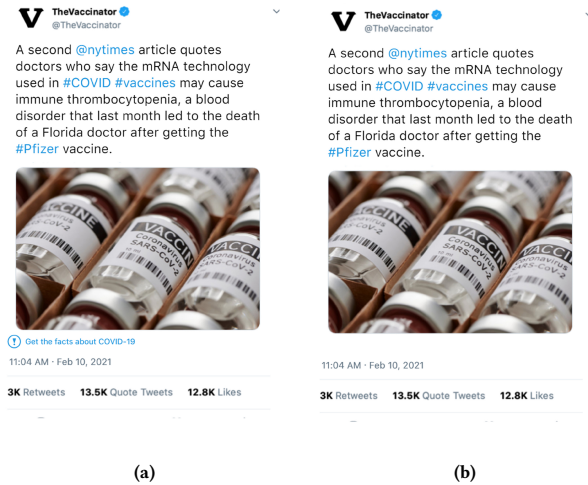


Figure 1: A Misleading Tweet: (a) With a Contextual Tag; (b) Without a Contextual Tag for Misleading Information.

To test the first hypothesis we utilized the Tweets containing *misleading information* shown in Figure 1a and Figure 1b. The Tweet in Figure 1a shows a warning-contextual tag underneath a Tweet, indicating that the content is labeled as misinformation. The Tweet promulgates COVID-19 misinformation about a rare adverse effect that was linked to the SARS-CoV-2 virus, not the vaccine, at the time of writing [5]. To remove bias due to the “influencer” effect, the Tweet comes from a verified account named “TheVaccinator” (which we made up) and indicates a relatively high interaction engagement with 3k retweets, 13.5k quotations, and 12.8k likes, which is consistent with the expected engagement of Tweets containing COVID-19 vaccine information [52]. An alteration of the same Tweet is shown in Figure 1b without the accompanying warning-contextual tag. To test the second hypothesis we utilized the Tweets containing *misleading information* shown in Figure 1b and Figure 2



Figure 4: **A Warning-An Interstitial** Cover Preceding a **Misleading Verified** Tweet



Figure 2: **An Interstitial** Cover Preceding a **Misleading** Tweet

(which includes **a warning-an interstitial** cover instead of a **warning contextual** tag).



(a)

(b)

Figure 3: **A Misleading Verified** Tweet: (a) **Without** a **Warning-Contextual** Tag; (b) **Without** a **Warning-Contextual** Tag for **Misleading** Information.

To test the third hypothesis we utilized the Tweets containing *verified information* shown in Figure 3a and Figure 3b. The Tweet content indicates the verified information distributed by the CDC about proceeding with the second dose of the COVID-19 vaccine in case an individual has a serious reaction from the first dose, altered to include a warning tag in Figure 3b [16]. To control for bias, the Tweet comes from a verified account “TheVirusMonitor” instead of the CDC account and indicates a similar engagement as the misleading Tweet [52]. To test the fourth hypothesis we utilized the Tweets containing *verified information* shown in Figure 3a and Figure 4. We retained Figure 3a for the comparison of the conditions and altered the **labeling-warnings** in the Figure 3b to include **a warning-an interstitial** cover instead of a warning tag.

A Verified Tweet: (a) Without a Warning Tag; (b) With a Warning Tag for Misleading Information.

A Warning Cover Preceding a Verified Tweet

3.2 Belief Echoes: Safety and Herd Immunity

Assuming the preconditions of the belief echoes are met, we examined the relationship between COVID-19 vaccine beliefs on safety and herd immunity and the perceived accuracy of Tweets with COVID-19 vaccine information in presence/absence of **warning labels-warnings**. We used the same Tweets from Figures 1-4 together to test the following hypotheses:

- H5a: The belief that COVID-19 vaccines are not safe will not affect the perception of accuracy of a Tweet with misleading information about COVID-19 in any condition (with a **warning-contextual** tag/**interstitial** cover or without any warning)
- H5b: The belief that COVID-19 vaccines are not safe will not affect the perception of accuracy of a Tweet with verified information about COVID-19 in any condition (with a **warning-contextual** tag/**interstitial** cover or without any warning)
- H6a: The belief that there is no need for COVID-19 vaccine because herd immunity exists will not affect the perception of accuracy of a Tweet with misleading information about COVID-19 in any condition (with a **warning-contextual** tag/**interstitial** cover or without any warning)
- H6b: The belief that there is no need for COVID-19 vaccine because herd immunity exists will not affect the perception of accuracy of a Tweet with verified information about COVID-19 in any condition (with a **warning-contextual** tag/**interstitial** cover or without any warning)

3.3 Belief Echoes: Efficacy and Hesitancy

Next, we examined the relationship between COVID-19 vaccine efficacy/hesitancy and the perceived accuracy of Tweets with COVID-19 vaccine information in presence/absence of warning [labels](#)[tags](#) [and covers](#). We used the same Tweets as shown in Figures 1-4 together to test the following hypotheses:

- H7a: The perception of producing efficacious COVID-19 vaccine will not affect the perception of accuracy of a Tweet with misleading information about COVID-19 in any condition (with a [warning-contextual](#) tag/[interstitial](#) cover or without any warning)
- H7b: The perception of producing efficacious COVID-19 vaccine will not affect the perception of accuracy of a Tweet with verified information about COVID-19 in any condition (with a [warning-contextual](#) tag/[interstitial](#) cover or without any warning)
- H8a: The COVID-19 vaccine personal hesitancy will not affect the perception of accuracy of a Tweet with misleading information about COVID-19 in any condition (with a [warning-contextual](#) tag/[interstitial](#) cover or without any warning)
- H8b: The COVID-19 vaccine personal hesitancy will not affect the perception of accuracy of a Tweet with verified information about COVID-19 in any condition (with a [warning-contextual](#) tag/[interstitial](#) cover or without any warning)
- H9a: The COVID-19 vaccine hesitancy for children will not affect the perception of accuracy of a Tweet with misleading information about COVID-19 in any condition (with a [warning-contextual](#) tag/[interstitial](#) cover or without any warning)
- H9b: The COVID-19 vaccine hesitancy for children will not affect the perception of accuracy of a Tweet with verified information about COVID-19 in any condition (with a [warning-contextual](#) tag/[interstitial](#) cover or without any warning)

3.4 Belief Echoes and Political Affiliation

To test the association between one's political affiliation and the perceived accuracy of the Tweets from Figures 1-4, following the evidence in [52] about the interplay between political affiliation and misinformation [warning-labels](#)[warnings](#), we asked:

- RQ1: Is there a difference in the perceived accuracy of COVID-19 misleading/verified Tweets with [warning-labels](#) [warnings](#) (tags or covers) between Republican, Democrat, and Independent users?
- RQ2: Is there a difference between the beliefs and subjective attitudes of the Twitter users about the COVID-19 vaccine based on their political affiliation?

3.5 Belief Echoes and Demographics

[To test the association between one's age, gender, education and the perceived accuracy of the Tweets from Figures 1-4, following the evidence in \[52\] about the interplay between political affiliation and misinformation warnings, we asked:](#)

- [RQ3a: Is there a difference in the perceived accuracy of COVID-19 misleading/verified Tweets with warnings \(tags or covers\) between users of different ages?](#)
- [RQ3b: Is there a difference between the beliefs and subjective attitudes of the Twitter users about the COVID-19 vaccine based on their age?](#)
- [RQ4a: Is there a difference in the perceived accuracy of COVID-19 misleading/verified Tweets with warnings \(tags or covers\) between users of different gender identity?](#)
- [RQ4b: Is there a difference between the beliefs and subjective attitudes of the Twitter users about the COVID-19 vaccine based on their gender identity?](#)
- [RQ5a: Is there a difference in the perceived accuracy of COVID-19 misleading/verified Tweets with warnings \(tags or covers\) between users of different education levels?](#)
- [RQ6b: Is there a difference between the beliefs and subjective attitudes of the Twitter users about the COVID-19 vaccine based on their education level?](#)

3.6 Sampling and Instrumentation

We first got approval from our Institutional Review Board (IRB) for an anonymous, non-full disclosure study. We set [out](#) to sample a population of US residents using Amazon Mechanical Turk that [is were](#) 18 years or above old, [is a Twitter user, and has Twitter users, and have](#) encountered at least one Tweet into their [feed that relates feeds that related](#) to COVID-19 vaccines. There were both reputation and attention checks to prevent [input](#) from bots and poor responses. The survey took between 5 and 10 minutes and the participants were compensated with the standard rate for participation. The study [questionnaire survey](#), incorporating the instruments from [2, 7], is provided in the Appendix. We utilized a 2x3 between group experimental design where participants were randomized into one of six groups: (1) misleading Tweet with a [warning-contextual](#) tag; (2) misleading Tweet without a [warning-contextual](#) tag; (3) misleading Tweet with [a warning-an interstitial](#) cover; (4) verified Tweet; (5) verified Tweet with a [warning-contextual](#) tag; and (6) verified Tweet with [a warning-an interstitial](#) cover.

After participation, the participants were debriefed and offered the option to revoke their answers. We crafted the content of the Tweets to be of relevance to the participants such that they [could](#) meaningfully engage with the Tweets' content (i.e., their responses are not arbitrary). We assumed participants understood the Twitter interface and metrics and were aware of the COVID-19 pandemic in general. However, we acknowledge that the level of interest regarding the COVID-19 vaccines could vary among the individual participants, affecting the extent to which their responses reflect their opinions.

4 STUDY RESULTS

We conducted an online survey (N = 319) in January and February 2021. The power analysis conducted with G* Power 3.1 [13] revealed that our sample was large enough to yield valid results for both Wilcoxon-Mann-Whitney U-test comparing two groups, [Kruskal-Wallis test comparing three or more groups](#), and Pearson's correlation (minimum of 44 per group). There were 180 (56.4%) males and 133 (41.7%) females, with 6 participants (1.8%) identifying as trans males, non-binary or preferring not to answer. The age

Table 1: Preconditions Tests: Hypotheses H1 to H4

	U-test	Significance
H1	$U = 1620.5$	$p = .217$
H2	$U = 1841$	$p = .004^*$
H3	$U = 1124$	$p = .063$
H4	$U = 1123.5$	$p = .087$
Significance Level: $\alpha = 0.05$		

brackets in the sample were distributed as follows: 13 (4.1%) [18 - 24], 108 (33.9%) [25 - 34], 98 (30.1%) [35 - 44], 46 (14.1%) [45 - 54], 37 (11.6%) [55 - 64], 17 (5.3%) [65 - 74], and 2 (.6%) [75 - 84]. In terms of education, 4 (1.3%) had less than high school, 24 (7.6%) had a high school degree or equivalent, 54 (17.0%) had some college but no degree, 43 (13.6%) had a 2-year degree, 139 (43.8%) had a 4-year college degree, and 53 (16.7%) had a graduate or professional degree. Our sample, while balanced on the other demographics, was Democrat-leaning with 59 (18.5%) Republicans, 157 (49.2%) Democrats, and 102 (32%) Independent participants. [The descriptive statistics of the variables used in the survey are given in Table 10 and Table 11 in the Appendix.](#)

4.1 Belief Echoes: Preconditions

[The To](#) test of preconditions for formation of belief echoes on Twitter about the COVID-19 vaccine, we first hypothesized that the presence of a [warning-contextual](#) tag under a Tweet containing misleading information about COVID-19 vaccines will not reduce the perceived accuracy of the Tweet's content relative to a no [warning-contextual](#) tag condition. [Indeed, the Mann-Whitney U test comparing the perceived accuracy between the participants exposed to the Tweet from We used a Wilcoxon-Mann-Whitney U-test to compare whether there is a difference in the dependent variable \(1\)\(a\) perceived accuracy of the tweet in the Appendix between the groups exposed to Figure 1a and Figure 1b, respectively. The test](#) was insignificant, as shown in Table 1. Confirming the H1 hypothesis, the results suggest that the [warning tags didn't have the intended effect on reducing the contextual tags did not reduce the](#) perceived accuracy of a misleading COVID-19 vaccine information on Twitter. This proves the existence of preconditions for echoing one's belief irrespective of the [warning-contextual](#) tags.

However, this [is was](#) not the case with the [warning covers used to label interstitial covers labeling](#) misleading information. The [Mann-Whitney U-test comparing the perceived accuracy Wilcoxon-Mann-Whitney U-test comparing whether there is a difference in the dependent variable \(1\)\(a\) perceived accuracy of the tweet](#) between the participants exposed to the Tweet from Figure 1b and Figure 2 was significant, rejecting the H2 hypothesis. The participants exposed to the [warning-interstitial](#) cover (Figure 2) reported, on average, that the Tweet was "not very accurate" while the ones exposed only to the Tweet content (Figure 1b) that the Tweet was "somewhat accurate." The [warning-interstitial](#) covers showed the intended effect of decreasing the perceived accuracy of misleading COVID-19 vaccine information on Twitter, dispelling one's echo beliefs [for Tweets labeled with covers in this case.](#)

As we suspected, the [warning labels didn't reduce the verified warnings did not reduce the verified](#) COVID-19 information, regardless of whether [a warning cover or warning an interstitial cover](#)

[or a contextual](#) tag was presented as a soft moderation intervention. The [Mann-Whitney U-tests comparing the perceived accuracy Wilcoxon-Mann-Whitney U-test comparing whether there is a difference in the dependent variable \(1\)\(a\) perceived accuracy of the tweet](#) between the participants exposed to the Tweet from Figure 3b and Figure 3a was insignificant as was the comparison between the exposures between Figure 3a and Figure 4 (Table 1). The retaining of H3 and H4 hypothesis [suggest suggests](#) that participants critically discern the content of a seemingly valid Tweet [, rather than the soft moderation labeling \(or alternatively, ignored the warnings altogether\).](#) Considering the previous reservations about [the soft moderation soft moderation tactics](#) implemented by Twitter when it comes to verified information [21], these results suggest that preconditions for echoing one's beliefs exist irrespective of any warning labeling also in the case where these beliefs are rooted in verifiable facts about COVID-19.

4.2 Belief Echoes: Safety and Herd Immunity

With the evidence of existing preconditions of belief echoes about [the COVID-19 vaccines](#), both for misleading and verified information, we set [out](#) to explore how these belief echoes materialize. We asked the participants to what extent [do](#) they agree with the following statement: "I am not favorable to the COVID-19 vaccines because I believe they are unsafe" [measured with the variable \(2\)\(a\) Beliefs: Safety.](#) We found negative correlation [with the perceived accuracy for the misleading between variables \(1\)\(a\) and \(2\)\(a\) for the misleading](#) Tweet with a [warning-contextual](#) tag (Figure 1a) and without a [warning-contextual](#) tag (Figure 1b) as shown in Table 2. The less participants were in favor of COVID-19 vaccines, the more accurate they perceived the misleading information regardless of the presence of a [warning-contextual](#) tag. We haven't found a significant correlation in the [warning-interstitial](#) cover condition.

[Safety and Perceived Accuracy Tests: H5a/b r-test Significance](#)
H5a — with a warning tag $r = -.612$ $p = .000^*$ **H5a** — without a warning tag $r = -.329$ $p = .017^*$ **H5b** — with a warning tag $r = -.344$ $p = .011^*$ **H5b** — with a warning cover $r = -.473$ $p = .000^*$

The test of the same relationship for the case of the [verified verified](#) Tweet also revealed a negative correlation [with the between the safety beliefs \(variable \(2\)\(a\)\) and the](#) Tweet's perceived accuracy (variable (1)(a)) in both the [warning-contextual](#) tag (Figure 3b) and the [warning-interstitial](#) cover condition (Figure 4). The more participants were in favor of COVID-19 vaccines the less accurate they perceived the verified information regardless if there was a [warning tag or a warning contextual tag or an interstitial](#) cover. On a first thought, this might be a surprising result, but a careful consideration indicates a presence of a belief echo and resistance to [soft moderated](#) verified information from [48], but not the standalone verified Twitter content.

Next, we asked the participants to what extent [do](#) they agree with the following statement: "There is no need to vaccinate for COVID-19 because I believe a natural herd immunity exists." [measured with the variable \(2\)\(b\) Beliefs: Herd Immunity.](#) We found negative correlation [with the between the beliefs on herd immunity \(variable \(2\)\(b\)\) and the](#) perceived accuracy for the [misleading Tweet with a warning misleading Tweet \(variable \(1\)\(a\)\) with a contextual](#)

Table 4: Efficacy and Perceived Accuracy Tests: H7a/b

	r-testU-test	Significance
H7a - with a warning contextual tag	$\chi^2(2) = 8.566$ $U = 8.566$	$p = .003^*$
H7a - with a warning interstitial cover	$\chi^2(2) = 9.237$ $U = 9.237$	$p = .002^*$
H6b - without a warning contextual tag	$\chi^2(2) = 3.969$ $U = 3.969$	$p = .005^*$
Significance Level: $\alpha = 0.05$		

Table 2: Safety and Perceived Accuracy Tests: H5a/b

	r-test	Significance
H5a - with a contextual tag	$r = -.612$	$p = .000^*$
H5a - without a contextual tag	$r = -.329$	$p = .017^*$
H5b - with a contextual tag	$r = -.344$	$p = .011^*$
H5b - with a contextual cover	$r = -.473$	$p = .000^*$
Significance Level: $\alpha = 0.05$		

tag (Figure 1a) and without a **warning contextual** tag (Figure 1b) as shown in Table 3. The more participants were in favor of the COVID-19 herd immunity, the more accurate they perceived the misleading information regardless if there was or **there wasn't a warning was not a contextual** tag. The test of the verified Tweet also revealed a negative correlation **with the between the beliefs on herd immunity (variable (2)(b)) and the** Tweet's perceived accuracy (variable (1)(a)) only in the original, no warning labels condition (Figure 3a). The more **the that** participants were in favor of COVID-19 herd immunity, the more accurate they perceived the verified information when no "soft moderation" intervention was applied. This finding adds to the evidence on the existence of **echo beliefs in the essential need of belief echoes that work against the** COVID-19 vaccines as a preferred way of **immunization gaining immunity**.

Table 3: Herd Immunity and Perceived Accuracy Tests: H6a/b

	r-test	Significance
H6a - with a warning contextual tag	$r = -.529$	$p = .000^*$
H6a - without a warning contextual tag	$r = -.387$	$p = .005^*$
H6b - without a warning contextual tag	$r = -.445$	$p = .001^*$
Significance Level: $\alpha = 0.05$		

4.3 Belief Echoes: Efficacy and Hesitancy

To test the subjective attitudes towards COVID-19 immunization, we asked the participants "will it be possible to produce efficacious COVID-19 vaccines". **We found a measured with the variable (3)(a) Efficacy. A Wilcoxon-Mann-Whitney U-test was used to compare**

whether there is a difference in the dependent variable (1)(a) Perceived accuracy of the tweet and the variable (3)(a) Efficacy and yielded a significant result for the misleading-misleading Tweet with a warning contextual tag (Figure 1a) and with a warning an interstitial cover (Figure 2) as shown in Table 4. In both cases, the participants who didn't did not believe in efficacious COVID-19 vaccines perceived the misleading Tweet "somewhat accurate" while the participants that did believe perceived it as "not very accurate." We also found a significant result for the verified-verified Tweet in its original form, without any "soft moderation" (Figure 3a). In this case, the participants that didn't believe in efficacious COVID-19 vaccines perceived the original, non-moderated Tweet "not very accurate." while the ones that did believe in the efficacy perceived it as "somewhat accurate."

To test the personal hesitancy to COVID-19 immunization, we asked the participants "Will you get vaccinated, if possible?" **We found a measured with the variable (3)(b) Personal hesitancy. A Kruskal-Wallis test was used to compare whether there is a difference in the dependent variable (1)(a) Perceived accuracy of the tweet and the variable (3)(b) Personal hesitancy and yielded a significant result for the misleading-misleading Tweet with a warning contextual tag (Figure 1a) and with a warning an interstitial cover (Figure 2) as shown in Figure-Table 5. In both cases, the participants that were hesitant to receive the COVID-19 vaccine perceived the misleading Tweet "somewhat accurate" while the participants that want to receive the vaccine as "not very accurate." The participants that were unsure, perceived the misleading Tweet in both cases as "not at all accurate." We also found a significant result for the verified-verified Tweet in its original form, without any "soft moderation" (Figure 3a). In the case, the participants that didn't want to receive the COVID-19 vaccine perceived the original, non-moderated Tweet "somewhat accurate." while the participants that wanted to receive the vaccine perceived it as "not very accurate." Similarly, the participants that were unsure perceived the Tweet as "not at all accurate."**

Table 5: Personal Hesitancy and Perceived Accuracy Tests: H8a/b

	r-test	Significance
H8a - with a warning contextual tag	$\chi^2(2) = 9.381$	$p = .009^*$
H8a - with a warning an interstitial cover	$\chi^2(2) = 7.163$	$p = .028^*$
H8b - without a warning contextual tag	$\chi^2(2) = 13.513$	$p = .001^*$
Significance Level: $\alpha = 0.05$		

To test the hesitancy to COVID-19 immunization for children, we asked the participants "Should children be vaccinated for COVID-19 too?" **We found a measured with the variable (3)(c) Hesitancy for children. A Wilcoxon-Mann-Whitney U-test was used to compare whether there is a difference in the dependent variable (1)(a) Perceived**

accuracy of the tweet and the variable (3)(c) Hesitancy for children and yielded a significant result for the misleading-misleading Tweet with only the warning tag (Figure 1a) as shown in Table 6. The participants that were hesitant to administer the COVID-19 vaccine to children perceived the misleading Tweet with a warning tag “somewhat accurate” while the participants that agreed with administering the COVID-19 vaccine to children “not very accurate.” We found a significant result for the verified-verified Tweet in its original form, without any “soft moderation” (Figure 3a). In the case, the participants that were hesitant to administer the COVID-19 to children perceived the an original, non-moderated Tweet as “somewhat accurate.” while the participants that agreed to administer the vaccine to children perceived it as “not very accurate.”

Table 6: Hesitancy for Children and Perceived Accuracy Tests: H9a/b

	r-test U-test	Significance
H9a - with a warning contextual tag	$\chi^2(2) = 10.663$ $U = 10.663$	$p = .001^*$
H9b - without a warning contextual tag	$\chi^2(2) = 8.001$ $U = 8.001$	$p = .005^*$
Significance Level: $\alpha = 0.05$		

4.4 Belief Echoes and Political Affiliation

Following the association between one’s political affiliation and the warnings of a misleading Twitter content [52], we analyzed the perceived accuracy among the participants based on their political affiliation (Republican, Democrat, independent). We found a Independent). A Kruskal-Wallis test was used to compare whether there is a difference in the dependent variable (1)(a) Perceived accuracy of the tweet and the variable (4)(d) Political leanings and yielded a significant difference in perception between the political affiliations of the participants for the misleading Tweet with the warning contextual tag and without the warning contextual tag as shown in Table 7. In both cases, the Republicans and independent participants perceived the Tweet as “somewhat accurate” while the Democrats as “not very accurate.”

Table 7: Political Affiliation and Perceived Accuracy Tests

	r-test	Significance
Misleading Tweet with a warning contextual tag	$\chi^2(2) = 7.063$	$p = .029^*$
Misleading Tweet without a warning contextual tag	$\chi^2(2) = 9.127$	$p = .005^*$
Significance Level: $\alpha = 0.05$		

We also analyzed the association between political affiliation, the beliefs of safety and herd immunity, and the subjective attitudes on the vaccine efficacy and hesitancy. While there were no

significant correlations about the safety and herd immunity beliefs, we the Pearson Chi-Square Test between dependent variables (3)(a) Efficacy and (3)(b) Personal hesitancy, and the variable (4)(d) Political leanings found significant differences on the question of producing efficacious vaccines, and personal hesitancy, respectively, as shown in Table 8. Almost one in four Republicans and and one in six Independents don’t expect to have an efficacious COVID-19 vaccine, while that proportion for the Democrats is one in forty. Half the Republicans and a third of the Independents are hesitant to receive the COVID-19 vaccine, while only a tenth of the Democrats won’t proceed with personal immunization. Roughly 40% of the Republicans and Independents are hesitant to vaccinate children for COVID-19, to which only 8.3% of the Democrats agree with.

Table 8: Political Affiliation, Beliefs, and Subjective Attitudes Tests

	r-test	Significance
Producing efficacious vaccines	$\chi(1) = 22.059$	$p = .001^*$
Personal Hesitancy	$\chi(2)(2) = 55.486$ $\chi(2) = 55.486$	$p = .000^*$
Hesitancy for Children	$\chi(1) = 45.665$	$p = .000^*$
Significance Level: $\alpha = 0.05$		

Belief Echoes and Demographics

Following the association between one’s demographic identity (age, gender, level of education) and the warnings of a misleading Twitter content [52], we analyzed the perceived accuracy among the participants based on their demographic groups using Kruskal-Wallis tests. We found no significant difference in perception when controlling for the age of the participants in any of the conditions shown in Figure 1-4 (RQ3a). Despite the lack of significance, we noticed that the participants below age of 45 perceived the verified Tweet without a contextual tag (Figure 3a) as “not very accurate” compared to the less sceptical participants of above 45 which perceived it as “somewhat accurate” ($\chi^2(6) = 9.580$ & $p = .0143$).

Controlling for the participant’s gender identity or level of education also did not yield any differences in perception (RQ4a). We noticed that the male participants perceived the verified Tweet with a contextual tag (Figure 4) as “somewhat accurate” while the female and gender non-binary participants as “not very accurate” ($\chi^2(2) = 3.339$ & $p = .0183$). Similarly, no statistically significant difference was noticed when controlled for level of education (RQ5a). A trend we noticed is the participants with lower level of education perceived the misleading Tweet with a contextual tag (Figure 1a) as “very accurate,” the ones with college degree as “not very accurate” and the ones with a graduate degree as “not at all accurate” ($\chi^2(5) = 8.741$ & $p = .0120$).

We also analyzed the association between the demographic categories, the beliefs of safety and herd immunity, and the subjective attitudes on the vaccine efficacy and hesitancy. While there were no significant correlations between the level of education and each belief (RQ5b), the Pearson Chi-Square Test between dependent variables (3)(a)

Efficacy and the variable (4)(a) Age and the Pearson Chi-Square Test between dependent variables (3)(a) Personal hesitancy and the variable (4)(b) Gender Identity, yielded statistically significant differences on the question of producing efficacious vaccines between the age groups (RQ3b) and on the question of the personal hesitancy and the gender identity of the participants (RQ4b) as shown in Table 9. Only 6.6% of the male participants and 0% of the non-binary participants did not expect to have an efficacious COVID-19 vaccine, however, that number was 15.5% for the female participants. Only 7.6% of the participants in the age bracket [18-24] are hesitant to receive the COVID-19 vaccine, but more than one in four of the other age groups won't proceed with personal immunization (25% [25-34]; 35.41% [35-44]; 28% [45-54]; 27% [55-64]; 29% [65-74]; 50% [75-above]).

Table 9: Demographics, Beliefs, and Subjective Attitudes Tests

	<u>r-test</u>	<u>Significance</u>
<u>Producing efficacious vaccines vs Age</u>	$\chi(6) = 31.566$	$p = .002^*$
<u>Personal Hesitancy vs Gender Identity</u>	$\chi(2) = 7.596$	$p = .023^*$
Significance Level: $\alpha = 0.05$		

5 DISCUSSION

Consistent with the previous evidence on receptivity to misinformation and resistance to warnings [7, 34], we found that the more likely participants were to believe that COVID-19 vaccines are unsafe, the more receptive they were to misleading COVID-19 vaccines information from Twitter, resisting the soft moderation intervention, proving the existence of belief echoes.

5.1 Strength and Type of Warning Label

That the participants perceived the misleading Tweets as accurate in the presence of a warning-contextual tag but not in the presence of a cover condition suggests that the warning-tags are not effective or insufficient to sway participants' perceived accuracy. These findings are consistent with previous research showing that the design of the warning-label-warnings affects individuals' perceptions of the content [7, 24, 31, 38]. The design of warning-labels-warnings and how they are presented to users can impact the warning-labels-warnings' effectiveness, with more explicit labels ones being more effective [31]. For instance in the previous study we build upon, it was found that individuals routinely ignored contextual warnings, which were akin to the warning-tags in our study in that they did tags because the tags do not obscure the misleading content nor require individuals to click through to see the information [24].

However, interstitial warnings, akin to cover warnings in our study in that they required the interstitial covers do require individuals to click through to continue, and as such, were effective in countering disinformation. Authors in [24] posit the COVID-19 vaccine misinformation. We concur that this may be because interstitial designs are more noticeable for users and thereby more

effective at countering misinformation since they present an actual "friction" for the user to engage with the content of their particular interest [9]. It may also be that these designs require users' engagement and thus necessitate a cognitive awareness of the tags-cover's textual content. Similarly, authors in [7] found that the perceived accuracy decreased with the increasing strength of the warning-contextual tags, such that tags which said "rated false" were significantly more effective than "disputed" tags at reducing beliefs in the misleading information. These findings suggest that more explicit, unambiguous warnings are more effective at countering misleading information. In our case, the warning or on Twitter originally, the interstitial covers - not the tags, which are more verbose and exact, were a more potent way of urging users to critically discern COVID-19 vaccine content on Twitter. However, users may habituate to the cover warnings in the long-term if they perceive that the moderator, Twitter, is biased in labeling content from users with particular political leanings [4].

Preconceived Notions-Explaining Notions - Explaining Results of Verified Tweets

Dispelling belief echoes on Twitter might be a more altogether on Twitter is most likely a complex task, since it is dependent on the content or type of misinformation and the subjective involvement of the participants. The fact that participants who were more likely to believe that COVID-19 vaccines were unsafe, were less likely to perceive the verified Tweet as inaccurate in both the warning-tag and-contextual tag and interstitial cover conditions may be due to the fact that the verified Tweet, though accurate, still reflected information that was negative about COVID-19 vaccines (i.e., invoking the idea that they may lead to serious side-effects and should be avoided by some participants in some situations). In other words, it would make sense that participants favouring vaccines would be more likely to disbelieve the Tweet expressing a concern about COVID-19 vaccines when accompanied by a warning label. A similar conclusion could be drawn also in the case of belief in herd immunity versus mass immunization with COVID-19 vaccines.

5.2 COVID-19 Vaccine Beliefs

In terms of efficacy, participants who thought that the COVID-19 vaccines were ineffective were more likely to rate the misleading Tweets as accurate, regardless of the soft moderation applied. This suggests that belief echoes persist despite the warnings, and perhaps, a presence of a warning-tag or-contextual tag or an interstitial cover may actually increase people's likelihood of finding a misleading Tweet accurate if it conforms to existing beliefs. This finding is consistent with the backfire effect previously observed for polarizing content on social media [36]. For example, evidence suggests that corrections on misleading Tweets strengthened misperceptions (or perceptions of accuracy) among those most strongly committed to the belief [34]. The corrections that contradict users' preconceived notions were found to lead individuals to double down on their beliefs.

In terms of hesitancy (both personal and for children vaccination), we found a similar effect, such that those who were hesitant about vaccines were more likely to perceive misleading Tweets with tags-and-contextual tags and interstitial covers as accurate while

those who wanted the vaccine perceived it as inaccurate. Again, the fact that only Tweets with tags and covers were viewed as accurate suggests evidence for a backfire effect such that the mere presence of the ~~tags/covers/warnings~~ may increase individuals' beliefs in the Tweets' accuracy if the content reinforces the participants' anti-vax stance. A similar conclusion follows from the fact that, for the original (verified) Tweet, the pro-vax participants who wanted to get the vaccine viewed ~~it-the content~~ as not very accurate but the anti-vax participants viewed it as accurate.

5.3 Political Affiliations

Along the lines of the findings in [6, 34, 52], we found further evidence of the association between user's political affiliations and the receptivity to misleading content. The Republican and Independent participants perceived the Tweet as "somewhat accurate" while the Democrat participants perceived it as "not very accurate" in both the misleading Tweet with and without a warning tag. That the difference between the expectation of an efficacious COVID-19 vaccine is ~~twentyfold-between-the~~ twenty-fold between Republicans and Democrats is a bit surprising, but consistent with the breakdown of trust in scientists to deliver an efficacious COVID-19 vaccine along the party line [18]. The hesitancy we found in our study is consistent with the previous reported breakdown for the COVID-19 vaccine ~~hesitant~~ hesitancy in Republicans and Democrats, both personally and in regards to children's vaccination [26]. Interestingly, Independents showed a high hesitancy on par with ~~the~~ Republicans in both cases.

Authors in [6] found that while corrections had the intended effect among Democrats, soft moderation techniques backfired among Republicans. Specifically, the authors found that while corrections of misleading claims decreased Democrats' perceptions of claim accuracy, they actually strengthened Republicans' perceptions of accuracy. As in [34], these findings suggest that corrections of misleading information on social media may not only be ineffective among some individuals but may actually reinforce individuals' preconceived notions. While our study did not assess participants' beliefs before and after receiving corrections, as all participants were only assessed once, the findings that political affiliation affects individuals' perceptions of accuracy and the impact that ~~warning labels~~ warnings have on those perceptions are consistent with the backfiring effect among individuals with certain political ideologies.

5.4 Demographics

~~Reluctance to heed security or privacy warnings is not a new phenomenon and has been well researched in the past [12, 19, 33]. While efforts have been invested in increasing the clarity of~~ Though without a statistical significance, the interplay between the demographics and the messages and design of affordances to attract attention and motivate users, habituation is a complex problem transcending security designs. Habituation describes a diminished response with repetitions of the same stimulus, decreasing the intended effect of security and privacy warnings among users. Authors in [49], in this context, have uncovered the phenomenon of "generalization" where habituation to one stimulus carries over to other novel stimuli that are similar in appearance. We didn't explore the diminished response with repetitions of perception of sheds a light on how people interact with Twitter content subjected to soft moderation.

The level of education, as expected, shows a correlation trend with the perception and interpretation of content under warning. Interestingly, the female and non-binary participants were more inclined to skepticism, e.g. to heed a contextual tag than their male counterparts even if the tag is applied to otherwise verified content. Skepticism seems to be inversely related to age given that participants below the age of 45 saw the verified Tweet as rather inaccurate. It is therefore worthy in the ~~same warning label to a Tweet, being that a tag or a cover, but that certainly warrants close research attention:~~ future to further explore the level of skepticism associated with or steaming from the belief echoes when controlled for demographic information as fitting into the broader trend of "generalised scepticism" when people navigate information on social media [15]. In our case, we are particularly interested about situations where users are given preference to choose between contextual indicators (covers) or interstitial indicators (tags) and how that choice affects the manifestation of belief echoes (in our current study, we followed the current scenario where Twitter instead chooses what label is presented based on the content and not on the user preference).

Further evidence in the context of skepticism gives our demographic analysis of the beliefs of safety and immunity. That females are more than twice as skeptical of an efficacious COVID-19 vaccine does not come to a surprise given the particular history of vaccine misinformation going back to the rumors of autism threat to newborns [41]. What is surprising to see is that the level of education is not a factor in deciding to immunize against COVID-19 and that hesitancy is rampant among the adults. Granted, we conducted our research during a period when the vaccine was scarce and surrounded by a cloud of rumours and doubt about the potential averse effects. Future research, therefore, should re-assess if this indeed was a result of the particular conditions of vaccine scarcity and adults indeed got their vaccine later [18] or the skepticism that persists in a spiral-of-silence [42] or privacy unravelling form [51]

~

5.5 Usable Security and Privacy Implications

Reluctance to heed security or privacy warnings is not a new phenomenon and has been well researched in the past [12, 19, 33]. The findings of our study suggest that heeding a misleading information warning only happened when the information is obscured by a plain text warning of the risks, not when the warning follows the Tweet with tag.

~~The warning~~ The contextual tag, consisting of an exclamation mark symbol urges users to "Get the facts about COVID-19," in Twitter's blue font, communicates a seemingly ambiguous message without explicitly addressing that the Tweet's content aims to mislead users about the COVID-19 vaccine. ~~Perhaps a line of research could explore a variation of~~ [35]. We therefore would propose a variation of contextual tags that are more direct, for example "This is COVID-19 misinformation", written in bold red font and conventional warning favicons. ~~Alternatively an impartial message like "No judgment, but this might be COVID-19 misinformation" could also show users' receptivity to not-so-overt moderation. The warning~~ The interstitial cover, along these lines, communicated a verbose message where Twitter appeared not taking sides by saying ~~~This Tweet violated the Twitter Rules about spreading misleading~~

~~and potentially harmful information related to COVID-19. Twitter has determined that may be in the public interest for the Tweet to Remain Accessible. It also provided with a link for the participant user to Learn More, which largely subsumes the warning contextual tag by leading users to a repository of verified COVID-19 information. Alternative wording like This Tweet was rated 'false', but we keep it in the public interest, based on the previous evidence [7], could yield a stronger reduction of misperception.~~

In the context of generalization, Twitter labels alternative narratives that are: (1) statements or assertions that have been confirmed to be false or misleading by subject-matter experts, such as public health authorities (~~misleading information~~); (2) statements or assertions in which the accuracy, truthfulness, or credibility of the claim is contested or unknown (~~disputed claims~~); and (3) information (which could be true or false) that is unconfirmed at the time it is shared (~~unverified claims~~) [37]. While efforts have been invested in increasing the clarity of the messages and design of affordances to attract attention and motivate users, habituation is a complex problem transcending security designs. Habituation describes a diminished response with repetitions of the same stimulus, decreasing the intended effect of security and privacy warnings among users [49]. With a similarity in labeling alternative narratives, from a user perspective, habituation to a tag or a cover for misleading COVID-19 vaccination could potentially carry over to other warnings about other polarizing events, such as elections. ~~An interesting line of research could investigate the generalization effect not just between different labels, but in various combinations of formatting and wording.~~ A user might be well aware and agree that some claims about elections are disputed, but they can nonetheless retain their beliefs about the COVID-19 vaccine safety and efficacy. Further so, ~~another~~ a line of research could be an investigation of the ~~generalization phenomenon~~ habituation between social media platforms, for example between warning labels on Twitter and Facebook.

~~This discussion brings an important aspects.~~ Our results highlight an important aspect of usable security affordances that ~~depart~~ departs from the conventional warning about system-level exploits ~~towards~~ toward content-level warnings. System-level exploits hardly relate to any potent beliefs (outside perhaps of the stereotypical foreign nation-state interference) or better said, users might not have strong polarizing stances on phishing or malware, usually perceiving it as a "bad thing" [14]. Content-level exploits, on the other side, are far more complex and potent in polarizing users, given that they are subjectively involved with the content [45][24, 45]. Users might ignore a red screen proceeding to a suspicious website, but they usually trust Chrome or Firefox that they have honest intentions in warning them about potential risks. Evidence already indicates that users are not trusting of ~~the~~ soft moderation intervention, feeling that Twitter itself ~~was is~~ biased and mislabeling content [21]. Remaining impartial while trying to dispel belief echoes might be a harder problem depending on the content - while there are safe and unsafe websites, there is ~~and will continue to be~~, a wealth of polarizing content on Twitter that will require content-relevant ~~warning labeling~~ warnings.

5.6 Ethical Implications

While we set out to investigate the effect of soft moderation on Twitter and debriefed the participants at the end of the study, the

results could have several ethical implications, nonetheless. We exposed the participants to a misleading and manipulated soft moderation of twitter content on the topic of the COVID-19 vaccine that could potentially affect participants' stance on vaccination and the pandemic. The exposure might not sway participants on the hesitancy or their perceptions of safety and efficacy, but could make the participants reconsider their approach of obtaining the vaccine for themselves or their children. The exposure could also affect the participants' stance of social media soft moderation in general and nudge people to move to ~~less regulated platforms~~ alt-platforms [53]. A recent example of such a migration from Twitter to Parler, Rumble and Newsmax was witnessed after Twitter actively labeled and removed false information on the platform during the 2020 U.S. elections [22].

That the participants were able to critically discern the content of the verified Tweet despite our ~~alternation to include warning labels~~ alteration to include warnings is reassuring and suggests that misinformation has the potential to be contained, if not eradicated, from social media platforms. However, the potential of crafting software that could silently attach or remove warning covers before they are presented to Twitter users could have unintended consequences. In the past, such an effort was tested in manipulating a Twitter textual content (not any additional affordances in the user interface) to induce misperceptions about the relationship between vaccines and autism [41]. With the evidence of nation-states censoring Twitter regarding narratives countering their interest in the past, it is possible that such a nation-state could use a similar approach and implement a "post-soft moderation" logic within a state-approved and disseminated social media application [47]. This may be far from the realm of possibility, even if the capabilities exist, but for such a sensitive topic as COVID-19 vaccination, meddling with the ~~warning labels~~ warnings could give an edge to a vaccine competitor in the global race for development and procurement of COVID-19 vaccines. We condemn such ideas and use of our research results. Evidence for such a nefarious misinformation Twitter campaign that promotes a homegrown Russian vaccine and undercuts rivals has already surfaced [17].

Perhaps outside the scope of this study, the ethical questions remains whether Twitter, or any social media platform, acting as a private entity, could set a precedent of an ultimate arbiter of what constitutes misinformation and what does not. Twitter most likely applies an automated means of warning labeling in conjunction with manual moderation [23], as evidence with the strange labeling of Tweets that contained the words "oxygen" and "frequency" for COVID-19 related Tweets [52]. Even with an attempt at honest moderation, cross-checked with the health authorities like CDC, a potential problem might arise in case a previously held belief, or a fact about COVID-19 ~~is being~~ later disputed. For example, at the beginning of the pandemic, authorities claimed masks were not effective in protecting ~~the virus from against the virus~~ spreading, a claim that later was not reversed, resulting in masks becoming essential to any human-to-human interaction [54]. So if the warning labels were applied to moderate any Tweet that contains the words "mask" and "stop" or "spread" at the early periods of the pandemic, they must be retracted. Such a thing could cast doubt on studies like ours, even if we as researchers, and Twitter as moderators, acted in good faith. Certainly, this could damage the reputation of users

as well as Twitter and further exacerbate the impression of not-so-honest impartiality in labeling content as misleading, especially against users identifying themselves as conservatives [4].

Future Research Generalization of Results

We acknowledge that there is further research to be done into investigating the full ramifications of soft moderation by social media platforms, especially beyond the topics of the COVID-19 pandemic or presidential elections. A promising line of research is the combination of soft and hard moderation, given that Twitter has exercised the right to ban or suspend accounts indefinitely that have been labeled for misinformation in the past, like in case of Donald Trump. One could probe the warning labeling algorithm and reverse engineer it to find if there is a relationship between the number/type of warning labels an account could receive, for example, before it gets permanently banned (if an automated ban exists, given the wide latitude and the imperative of Twitter to remain not overly controlling of the public discourse). When social media platforms, as private entities, are predetermined by users to hold biased positions [21], any action taken to apply soft moderation techniques may be undermined in the process, instead working to legitimize the beliefs of skeptics. In this direction, a content analysis of the Tweets being labeled could reveal the topics, images, words, or the network of accounts behind such impressions.

In the United States, where the right to speech is protected, more research may be done to see how alternative narratives, belonging to a same type of content or a topic used a couple of Tweets that were relevant to the state of the pandemic and mass immunization during the period of January-February 2021, which could be perceived with a different level of accuracy after a certain period of time. The particular choice of the Tweet's content as well as the engagement metrics might affect how users perceived the Tweets even though we attempted to remain impartial as possible by selecting content from the NPR for the Tweet in Figure 1 and 2 [5] and CDC for the Tweet in Figure 3 and 4 [16]. We opted for increased engagement with the Tweets following the findings in [52] relevant for Tweets that have been moderated by Twitter with interstitial covers and/or contextual tags. Because we did not vary the engagement metrics and content we cannot exclude the possibility that the particular selections might induce some participants to act based on the Tweet's content credibility (e.g. "This is from the New York Times or CDC so it must be a verified information") or engagement credibility (e.g. COVID-19 vaccines cause adverse effects leading to death) are soft moderated between platforms, for example comparing Twitter, Facebook, or Parler. Soft moderated content is usually closely related to content used for trolling so this relationship could be also explored, such as understanding if warning labeled Tweets provoke emotional response and of what kind. Similarly, the warning labeling can be associated to identify the evolution of political information operations on Twitter, that have been waged on the topic of COVID-19 already [46]: "The numbers are high so this information is probably valid". Other confounding factors such as nuanced political stances, religious beliefs, or occupational hazards might have had implicit effect unaccounted for during the hypotheses testing in our study.

Scope Limitations

The current study has important limitations. First, it is possible that a different topic or even different information regarding the effect of the COVID-19 vaccines would have different outcomes. We used a couple of Tweets that were relevant to the state of the pandemic and mass immunization during the period of January-February 2021, which could be perceived with a different level of accuracy after a certain period of time. We used only two Tweets and a study that explores the effect on multiple misleading or verified Tweets on COVID-19 vaccine could uncover different efficiency or strength of soft moderation. Overall, the findings in the present study may be specific to the effect that warning tags have on COVID-19 (mis)information and cannot be generalized to other topics. Second, participants who are more regular social media users in general may be desensitized to the information presented in the Tweets, which may have affected their perception of the issue irrespective of the warnings. Using a small number of stimuli is a limitation, stemming both from restricted financial resources and limited attention span of participants, also pertaining the general study of misinformation warnings on which we build upon [24]. To address this limitation and gather more evidence that the interstitial warnings work but the contextual tags do not, we replicated the study with the same stimuli in smart home settings where Amazon Alexa was the one that read the Tweets and the warnings to the participants [39]. Here too, users heed the interstitial warnings spoken back by Alexa before and ignored the contextual tags. Even more, with an audio instead of a visual interface, Alexa was able to "convince" the participant to perceive the otherwise verified information (Figure 4) as misinformation. In another study we used content manipulation with adversarial hashtags and addition/removal of the contextual tags and found that the belief echoes affect how one perceives a tweet with a contextual tag [40]. The misinformation Tweet in this case included a rumour that President Biden is dropping Operation Warp Speed for the federal vaccination effort (the name was the only thing dropped, while keeping the operation in tact) [25]. Here, the vaccine hesitant participants deemed the Tweet as "somewhat accurate" despite the contextual tag.

Third, our experiment was limited to Twitter as a social media platform of choice. Since the content and images we present are borrowed and adapted to the study objectives from Twitter, we are limited to evaluating the effects of the warning tags and warning cover on perceptions of accuracy on Twitter only and may not and may not entirely be generalized regarding other social media platforms. We were limited to the formatting and wording of the warning labels chosen by Twitter at the time of the study. If Twitter chooses to place re-situate the tag, say by placing it on top of the Tweet instead of the bottom, the results could be different. Similarly, if the wording of the warning interstitial cover changes, the results might not hold for such new conditions. Fourth, we did not examine the effects over a period of time. Thus, we are unable to examine the Tweets' effects following the study. We also acknowledge another limitation imposed of the timeline of the study and the speed of vaccine development. We acknowledge that for future replication(s) of the study, any further developments in regards to the COVID-19 vaccine development. By the time participants completed the study, much more might be known about the COVID-19 vaccines to sway public opinion. For instance, if many participants have gotten the vaccine without major side effects by the time participants

complete the study, this might affect their responses. Fifth, although we tried to sample a representative set of participants for our study using Amazon Mechanical Turk, the outcomes might have been different if we used another platform, or other type of sampling. Also, a much vaccines in particular and the pandemic in general have to be incorporated in crafting the stimuli in order to remain in synchronization with relevant misinformation and Twitter's soft moderation policies. Future replication studies would certainly benefit if the effect of the Twitter warnings is tested on a larger sample size could have provided in providing a more nuanced view of the soft moderation, but we had limited funding for this study general acceptance of soft moderation in online discourses.

6 CONCLUSION

In the present study, we sought to determine whether two forms of soft moderation - contextual covers and interstitial tags - on Twitter affect the perceived accuracy of Tweets pertaining to COVID-19 vaccines. We were also interested in examining whether perceived Tweet accuracy varies based on individual's beliefs regarding COVID-19 vaccine safety, efficacy, willingness to receive vaccinations, and their political affiliations. Overall, our results suggest that warning the interstitial covers are more effective than warning the contextual tags in dispelling individuals' beliefs about misleading Tweets, individuals' pre-existing beliefs regarding COVID-19 vaccine safety, efficacy, and hesitancy affect individuals' perceptions of Tweet accuracy such that their perceptions of the accuracy align with their biases. Furthermore, our results also show that individuals' political affiliations also affect their perceptions of accuracy for misleading Tweets such that Republicans and Independents, who are more likely to express skepticism regarding vaccines, are more likely to perceive misleading Tweets as accurate, irrespective of any moderation effort.

In all cases, individuals perceive the Tweets in ways that are most favourable to or consistent with their pre-existing beliefs. This may lead to a backfire effect, as evidenced by the fact that individuals who were skeptical of vaccines were more likely to rate misleading Tweets with tags and covers, but not misleading Tweets without tags, as accurate. Taken Skepticism for the information consumed on Twitter was also observed among the non-male Twitter users in regards the production of efficacious COVID-19 vaccines. We found that the adult Twitter in users are much more hesitant to get the COVID-19 vaccine than their younger (< 25 years) counterparts. Taken into consideration together, our findings provide additional evidence for the existence of belief echoes pertaining to COVID-19 vaccines that are largely resistant to soft moderation in the form of warning contextual tags but not warning interstitial covers. We believe that the insight gained from this research regarding how individuals' pre-existing belief biases impact perceptions of Tweet accuracy in the context of soft moderation can be used to develop more effective moderation techniques that do minimize unintended consequences. We hope that our results could inform the usable security community towards future steps in eradicating misinformation on Twitter and social media in general of exposure to misinformation.

REFERENCES

- [1] Jon-Patrick Allem and Emilio Ferrara. 2018. Could social bots pose a threat to public health? *American journal of public health* 108, 8 (2018), 1005.
- [2] Luigi Roberto Biasio, Guglielmo Bonaccorsi, Chiara Lorini, and Sergio Pecorelli. 2020. Assessing COVID-19 vaccine literacy: A preliminary online survey. *Human Vaccines & Immunotherapeutics* 0, 0 (2020), 1–9. <https://doi.org/10.1080/21645515.2020.1829315>
- [3] David A Broniatowski, Amelia M Jamison, SiHua Qi, Lulwah AlKulaib, Tao Chen, Adrian Benton, Sandra C Quinn, and Mark Dredze. 2018. Weaponized health communication: Twitter bots and Russian trolls amplify the vaccine debate. *American journal of public health* 108, 10 (2018), 1378–1384.
- [4] Jenna Burrell, Zoe Kahn, Anne Jonas, and Daniel Griffin. 2019. When Users Control the Algorithms: Values Expressed in Practices on Twitter. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 138 (Nov. 2019), 20 pages. <https://doi.org/10.1145/3359240>
- [5] Bill Chappell. 2021. Instagram Bars Robert F. Kennedy Jr. For Spreading Vaccine Misinformation. (2021). <https://www.npr.org/sections/coronavirus-live-updates/2021/02/11/966902737/instagram-bars-robert-f-kennedy-jr-for-spreading-vaccine-misinformation>.
- [6] Dino P Christenson, Sarah E Kreps, and Douglas L Kriner. 2020. Contemporary Presidency: Going Public in an Era of Social Media: Tweets, Corrections, and Public Opinion. *Presidential Studies Quarterly* (2020).
- [7] Katherine Clayton, Spencer Blair, Jonathan A Busam, Samuel Forstner, John Glance, Guy Green, Anna Kawata, Akhila Kovvuri, Jonathan Martin, Evan Morgan, et al. 2019. Real solutions for fake news? Measuring the effectiveness of general warnings and fact-check tags in reducing belief in false stories on social media. *Political Behavior* (2019), 1–23.
- [8] Warren Cornwall. 2020. Officials gird for a war on vaccine misinformation. *Science* 369, 6499 (2020), 14–15. <https://doi.org/10.1126/science.369.6499.14> arXiv:<https://science.sciencemag.org/content/369/6499/14.full.pdf>
- [9] Anna L. Cox, Sandy J.J. Gould, Marta E. Cecchinato, Ioanna Iacovides, and Ian Renfree. 2016. Design Frictions for Mindful Interactions: The Case for Microboundaries. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '16)*. Association for Computing Machinery, New York, NY, USA, 1389A–1397. <https://doi.org/10.1145/2851581.2892410>
- [10] Ulrich KH Ecker, Stephan Lewandowsky, and David TW Tang. 2010. Explicit warnings reduce but do not eliminate the continued influence of misinformation. *Memory & cognition* 38, 8 (2010), 1087–1100.
- [11] European Parliament. [n. d.]. Vaccine Hesitancy and Drop in Vaccination Rates in Europe - Thursday, 19 April 2018. https://www.europarl.europa.eu/doceo/document/TA-8-2018-0188_EN.html.
- [12] Michael Fagan and Mohammad Maifi Hasan Khan. 2016. Why Do They Do What They Do?: A Study of What Motivates Users to (Not) Follow Computer Security Advice. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*. USENIX Association, Denver, CO, 59–75. <https://www.usenix.org/conference/soups2016/technical-sessions/presentation/fagan>
- [13] Franz Faul, Edgar Erdfelder, Axel Buchner, and Albert-Georg Lang. 2009. Statistical power analyses using G*Power 3.1: Tests for correlation and regression analyses. *Behavior Research Methods* 41, 4 (2009), 1149–1160. <https://doi.org/10.3758/BRM.41.4.1149>
- [14] Adrienne Porter Felt, Alex Ainslie, Robert W Reeder, Sunny Consolvo, Somas Thyagaraja, Alan Bettis, Helen Harris, and Jeff Grimes. 2015. Improving SSL warnings: Comprehension and adherence. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems*. 2893–2902.
- [15] Richard Fletcher and Rasmus Kleis Nielsen. 2019. Generalised scepticism: how people navigate news on social media. *Information, Communication & Society* 22, 12 (2019), 1751–1769.
- [16] Centers for Disease Control and Prevention. 2021. COVID-19 Vaccines and Allergic Reactions. <https://www.cdc.gov/coronavirus/2019-ncov/vaccines/safety/allergic-reaction.html>.
- [17] Sheera Frenkel, Maria Abi-Habib, and Julian E. Barnes. 2021. Russian Campaign Promotes Homegrown Vaccine and Undercuts Rivals. <https://www.nytimes.com/2021/02/05/technology/russia-covid-vaccine-disinformation.html>.
- [18] Cary Funk and Alec Tyson. 2020. Intent to Get a COVID-19 Vaccine Rises to 60% as Confidence in Research and Development Process Increases. <https://www.pewresearch.org/science/2020/12/03/intent-to-get-a-covid-19-vaccine-rises-to-60-as-confidence-in-research-and-development-process>
- [19] Simson Garfinkel and Heather Richter Lipford. 2014. Usable security: History, themes, and challenges. *Synthesis Lectures on Information Security, Privacy, and Trust* 5, 2 (2014), 1–124.
- [20] Bertram Gawronski, Roland Deutsch, Sawan Mbirkou, Beate Seibt, and Fritz Strack. 2008. When “Just Say No” is not enough: Affirmation versus negation training and the reduction of automatic stereotype activation. *Journal of Experimental Social Psychology* 44, 2 (2008), 370–377. <https://doi.org/10.1016/j.jesp.2006.12.004>
- [21] Christine Geeng, Tiona Francisco, Jevin West, and Franziska Roesner. 2020. Social Media COVID-19 Misinformation Interventions Viewed Positively, But Have Limited Impact. arXiv:cs.CY/2012.11055
- [22] Mike Isaac and Kellen Browning. 2020. Fact-Checked on Facebook and Twitter, Conservatives Switch Their Apps. <https://www.nytimes.com/2020/11/11/technology/parler-rumble-newsmax.html>.

- [23] Peter Jachim, Filipo Sharevski, and Emma Pieroni. 2021. TrollHunter2020: Real-Time Detection of Trolling Narratives on Twitter During the 2020 U.S. Elections. In *Proceedings of the 2021 ACM Workshop on Security and Privacy Analytics (IWSPA '21)*. Association for Computing Machinery, New York, NY, USA, 55â\$65. <https://doi.org/10.1145/3445970.3451158>
- [24] Ben Kaiser, Jerry Wei, Eli Lucherini, Kevin Lee, J. Nathan Matias, and Jonathan Mayer. 2021. Adapting Security Warnings to Counter Online Disinformation. In *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, 1163–1180. <https://www.usenix.org/conference/usenixsecurity21/presentation/kaiser>
- [25] Kaplan, Sheila and Stolberg, Sheryl Gay. 2021. Biden picks a former F.D.A. chief to lead federal vaccine efforts, which will drop the name Operation Warp Speed. <https://www.nytimes.com/2021/01/15/world/biden-picks-a-former-fda-chief-to-lead-federal-vaccine-efforts-which-will-drop-the-name-operation-warp-speed.html>
- [26] Kendall Karson. 2020. Americans willing to receive COVID-19 vaccine but divided on timing: Poll. <https://abcnews.go.com/Politics/americans-receive-covid-19-vaccine-divided-timing-poll/story?id=74703426>
- [27] Jessica Keim-Malpass, Emma M. Mitchell, Emily Sun, and Christine Kennedy. 2017. Using Twitter to Understand Public Perceptions Regarding the #HPV Vaccine: Opportunities for Public Health Nurses to Engage in Social Marketing. *Public Health Nursing* 34, 4 (2017), 316–323. <https://doi.org/10.1111/phn.12318> <https://onlinelibrary.wiley.com/doi/pdf/10.1111/phn.12318>
- [28] Jeffrey V Lazarus, Scott C Ratzan, Adam Palayew, Lawrence O Gostin, Heidi J Larson, Kenneth Rabin, Spencer Kimball, and Ayman El-Mohandes. 2020. A global survey of potential acceptance of a COVID-19 vaccine. *Nature medicine* (2020), 1–4.
- [29] Joan Massachs, Corrado Monti, Gianmarco De Francisci Morales, and Francesco Bonchi. 2020. Roots of Trumpism: Homophily and Social Feedback in Donald Trump Support on Reddit. In *12th ACM Conference on Web Science*. 49–58.
- [30] Philip M Massey, Amy Leader, Elad Yom-Tov, Alexandra Budenz, Kara Fisher, and Ann C Klassen. 2016. Applying Multiple Data Collection Tools to Quantify Human Papillomavirus Vaccine Communication on Twitter. *J Med Internet Res* 18, 12 (05 Dec 2016), e318. <https://doi.org/10.2196/jmir.6670>
- [31] Patricia L. Moravec, Antino Kim, and Alan R. Dennis. 2020. Appealing to Sense and Sensibility: System 1 and System 2 Interventions for Fake News on Social Media. *Information Systems Research* 31, 3 (2020), 987–1006. <https://doi.org/10.1287/isre.2020.0927>
- [32] Adam Mosseri. 2016. Addressing Hoaxes and Fake News. *Facebook* (2016). <https://about.fb.com/news/2016/12/news-feed-fyi-addressing-hoaxes-and-fake-news/>
- [33] James Nicholson, Lynne Coventry, and Pam Briggs. 2017. Can we fight social engineering attacks by social means? Assessing social salience as a means to improve phishing detection. In *Thirteenth Symposium on Usable Privacy and Security (SOUPS 2017)*. USENIX Association, Santa Clara, CA, 285–298. <https://www.usenix.org/conference/soups2017/technical-sessions/presentation/nicholson>
- [34] Brendan Nyhan and Jason Reifler. 2010. When corrections fail: The persistence of political misperceptions. *Political Behavior* 32, 2 (2010), 303–330.
- [35] Anne Oeldorf-Hirsch, Mike Schmierbach, Alyssa Appelman, and Michael P. Boyle. 2020. The Ineffectiveness of Fact-Checking Labels on News Memes and Articles. *Mass Communication and Society* 23, 5 (2020), 682–704. <https://doi.org/10.1080/15205436.2020.1733613> [arXiv:https://doi.org/10.1080/15205436.2020.1733613](https://arxiv.org/abs/https://doi.org/10.1080/15205436.2020.1733613)
- [36] Gordon Pennycook, Adam Bear, Evan T Collins, and David G Rand. 2020. The implied truth effect: Attaching warnings to a subset of fake news headlines increases perceived accuracy of headlines without warnings. *Management Science* (2020).
- [37] Yoel Roth and Nick Pickles. 2020. Updating our approach to misleading information. *Twitter* (2020). https://blog.twitter.com/en_us/topics/product/2020/updating-our-approach-to-misleading-information.html
- [38] Haeseung Seo, Aiping Xiong, and Dongwon Lee. 2019. Trust It or Not: Effects of Machine-Learning Warnings in Helping Individuals Mitigate Misinformation. In *Proceedings of the 10th ACM Conference on Web Science (WebSci '19)*. Association for Computing Machinery, New York, NY, USA, 265â\$274. <https://doi.org/10.1145/3292522.3326012>
- [39] Filipo Sharevski and Donald Gover. 2021. Two Truths and a Lie: Exploring Soft Moderation of COVID-19 Misinformation with Amazon Alexa. In *The 16th International Conference on Availability, Reliability and Security (ARES 2021)*. Association for Computing Machinery, New York, NY, USA, Article 79, 9 pages. <https://doi.org/10.1145/3465481.3470017>
- [40] Filipo Sharevski, Alice Huff, Peter Jachim, and Emma Pieroni. 2021. (Mis)perceptions and Engagement on Twitter: COVID-19 Vaccine Rumors on Efficacy and Mass Immunization Effort. [arXiv:cs/2007.3869553](https://arxiv.org/abs/cs/2007.3869553)
- [41] Filipo Sharevski, Peter Jachim, and Kevin Florek. 2020. To Tweet or Not to Tweet: Covertly Manipulating a Twitter Debate on Vaccines Using Malware-Induced Misperceptions. In *Proceedings of the 15th International Conference on Availability, Reliability and Security (ARES '20)*. Association for Computing Machinery, New York, NY, USA, Article 75, 12 pages. <https://doi.org/10.1145/3407023.3407025>
- [42] Filipo Sharevski, Paige Treebridge, Peter Jachim, Audrey Li, Adam Babin, and Jessica Westbrook. 2020. Beyond Trolling: Malware-Induced Misperception Attacks on Polarized Facebook Discourse. [arXiv:cs.HC/2002.03885](https://arxiv.org/abs/cs.HC/2002.03885)
- [43] Jeff Smith. 2017. Designing Against Misinformation. *Medium* (2017). <https://medium.com/facebook-design/designing-against-misinformation-e5846b3aa1e2>
- [44] Todd Spangler. 2020. Twitter Flags 200 Trump Posts as False or Disputed Since Election Day - Variety. *Variety* (2020). <https://variety.com/2020/digital/news/twitter-trump-200-disputed-misleading-claims-election-1234841137/>
- [45] Leo G Stewart, Ahmer Arif, and Kate Starbird. 2018. Examining trolls and polarization with a retweet network. In *Proc. ACM WSDM, workshop on misinformation and misbehavior mining on the web*.
- [46] Benjamin Strick. 2020. Discovering A Pro-Chinese Government Information Operation On Twitter and Facebook: Analysis Of The #Miles-Guo Bot Network. <https://www.bellingcat.com/news/2020/05/05/uncovering-a-pro-chinese-government-information-operation-on-twitter-and-facebook-analysis-of-the-miles-guo-bot-network>
- [47] Kurt Thomas, Chris Grier, and Vern Paxson. 2012. Adapting social spam infrastructure for political censorship. In *5th {USENIX} Workshop on Large-Scale Exploits and Emergent Threats (LEET) 12*.
- [48] Emily Thorson. 2016. Belief echoes: The persistent effects of corrected misinformation. *Political Communication* 33, 3 (2016), 460–480.
- [49] Anthony Vance, David Eargle, Jeffrey L. Jenkins, C. Brock Kirwan, and Bonnie Brinton Anderson. 2019. The Fog of Warnings: How Non-essential Notifications Blur with Security Warnings. In *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*. USENIX Association, Santa Clara, CA. <https://www.usenix.org/conference/soups2019/presentation/vance>
- [50] Robin C. Vanderpool, Anna Gaysynsky, and Wen-Ying Sylvia Chou. 2020. Using a Global Pandemic as a Teachable Moment to Promote Vaccine Literacy and Build Resilience to Misinformation. *American Journal of Public Health* 110, S3 (2020), S284–S285. <https://doi.org/10.2105/AJPH.2020.305906>
- [51] Mark Warner, Andreas Gutmann, M. Angela Sasse, and Ann Blandford. 2018. Privacy Unraveling Around Explicit HIV Status Disclosure Fields in the Online Geosocial Hookup App Grindr. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW, Article 181 (Nov. 2018), 22 pages. <https://doi.org/10.1145/3274450>
- [52] Savvas Zannettou. 2021. "I Won the Election!": An Empirical Analysis of Soft Moderation Interventions on Twitter. [arXiv 2101.07183v1](https://arxiv.org/pdf/2101.07183v1) (18 January 2021). <https://arxiv.org/pdf/2101.07183v1>
- [53] Savvas Zannettou, Tristan Caulfield, Emiliano De Cristofaro, Nicolas Kourtellis, Ilias Leontiadis, Michael Sirivianos, Gianluca Stringhini, and Jeremy Blackburn. 2017. The Web Centipede: Understanding How Web Communities Influence Each Other through the Lens of Mainstream and Alternative News Sources. In *Proceedings of the 2017 Internet Measurement Conference (IMC '17)*. Association for Computing Machinery, New York, NY, USA, 405â\$417. <https://doi.org/10.1145/3131365.3131390>
- [54] Joyce C Zhang and Anil Adisesh. 2020. Controversies in respiratory protective equipment selection and use during COVID-19. *Journal of hospital medicine* 15, 5 (2020).

APPENDIX

The study questionnaire

Survey

The study survey included the following questions:

- (1) **Perceived Accuracy of a Tweet:** ~~4-~~
 - (a) *To the best of your knowledge, how accurate is the claim described in the Tweet?*
4-point Likert scale (1-not at all accurate, 2-not very accurate, 3-somewhat accurate, 4-very accurate).
- (2) **Beliefs:** ~~2-~~
 - (a) *How much do you agree with the following statement: "I am not favorable to vaccines because they are unsafe"?*
~~3-~~ 4-point Likert scale (1 - Totally, 2 - A Little, 3 - Partially, 4 - Not at All).
 - (b) *How much do you agree with the following statement: "here is no need to vaccinate because a natural immunity exists"?*
4-point Likert scale (1 - Totally, 2 - A Little, 3 - Partially, 4 - Not at All).
- (3) **Subjective Attitudes:** ~~4-~~

(a) *Will it be possible to produce safe and efficacious COVID-19 vaccines?*

Yes/No. ~~5-~~

(b) *Will you get vaccinated, if possible?*

Yes/No/I Don't Know. ~~6-~~

(c) *Should children be vaccinated for COVID-19 too?*

Yes/No.

(4) **Demographics:**

(a) age

(b) gender identity

(c) education

(d) political leanings

Descriptive Statistics

Table 10: Continuous Variables

<u>Variable</u>	<u>Mean</u>	<u>Std. Dev.</u>
Age, gender identity, education, political leanings-		
<u>(1)(a) - Figure 1a</u>	<u>2.13</u>	<u>.953</u>
<u>(1)(a) - Figure 1b</u>	<u>2.36</u>	<u>.942</u>
<u>(1)(a) - Figure 2</u>	<u>1.83</u>	<u>.753</u>
<u>(1)(a) - Figure 3a</u>	<u>2.35</u>	<u>.974</u>
<u>(1)(a) - Figure 3b</u>	<u>1.98</u>	<u>.939</u>
<u>(1)(a) - Figure 4</u>	<u>2.32</u>	<u>.956</u>
<u>(2)(a)</u>	<u>3.48</u>	<u>.918</u>
<u>(2)(b)</u>	<u>3.5</u>	<u>.909</u>

Table 11: Categorical Variables

<u>Variable</u>	<u>Yes</u>	<u>No</u>	<u>Don't Know</u>
<u>(3)(a)</u>	<u>89.7%</u>	<u>10.3%</u>	<u>N/A</u>
<u>(3)(b)</u>	<u>71.5%</u>	<u>21.9%</u>	<u>6.6%</u>
<u>(3)(c)</u>	<u>75.2%</u>	<u>24.8%</u>	<u>N/A</u>